

## AMENDMENTS TO THE CLAIMS

1        1. (Currently amended) A method for attempting to access a first data entity in  
2 a file system, the method being performed by a first computing entity, the file system  
3 also including one or more additional data entities that are concurrently accessible to at  
4 least one other computing entity, the file system including an owner field ~~for indicating if~~  
5 ~~that can be used to determine whether~~ the first data entity is leased by a computing  
6 entity and a time field ~~for indicating that can be used to determine~~ whether a lease for  
7 the first data entity has expired, the method comprising:

8              attempting to obtain a lease on the first data entity by ~~performing the~~  
9 ~~following substeps reading the owner field and:~~

10              ~~reserving a data storage unit containing the owner field;~~  
11              if the owner field indicates that the first data entity is not currently  
12 leased, writing to the owner field to indicate an assumption of a lease of  
13 the first data entity and writing to the time field to indicate when the lease  
14 expires; or

15              if the owner field indicates that the first data entity has been leased,  
16 ~~and the time field indicates that the lease is active, writing to the owner~~  
17 ~~field in a queue to indicate an intention to lease and writing to the time~~  
18 ~~field to indicate when the lease expires; and reading the time field and:~~

19              ~~if the owner field indicates that the first data entity has been~~  
20 ~~leased, but the time field indicates that the lease has expired,~~  
21 writing to the owner field to break the existing lease and to indicate  
22 an assumption of a new lease and writing to the time field to  
23 indicate when the new lease expires; and or

24              if the time field indicates that the lease is still active,  
25 concluding that the first data entity is currently unavailable; and

26              if a lease is obtained, accessing the first data entity while the lease is  
27 active in effect.

1        2. (New) The method of claim 1, wherein the first data entity is a file.

1           3. (New)   The method of claim 2, wherein the first data entity includes  
2 metadata and the owner field is located in this metadata.

1           4. (New)   The method of claim 1, wherein the first data entity is a directory.

1           5. (New)   The method of claim 1, wherein the step of writing to the owner field  
2 to indicate an assumption of a lease of the first data entity comprises writing a data  
3 value to the owner field that uniquely identifies the first computing entity.

1           6. (New)   The method of claim 5, wherein the data value that uniquely  
2 identifies the first computing entity is determined autonomously by the first computing  
3 entity.

1           7. (New)   The method of claim 5, wherein the owner field indicates that the  
2 first data entity is not currently leased when the owner field contains a value of zero.

1           8. (New)   The method of claim 1, wherein a lease expires a predetermined  
2 period of time after the lease begins, and wherein the step of writing to the time field to  
3 indicate when the lease expires comprises writing a current time value to the time field.

1           9. (New)   The method of claim 1, wherein the first computing entity  
2 determines whether a prior lease has expired by reading a first value from the time field,  
3 delaying for a predetermined lease period and reading a second value from the time  
4 field, wherein the first computing entity determines that the prior lease has expired if the  
5 second value is the same as the first value, and the first computing entity determines  
6 that the prior lease has not expired if the second value is different from the first value.

1           10. (New)   The method of claim 1, wherein the steps of reading the owner  
2 field and reading the time field are both performed in a single read operation.

1        11. (New)    The method of claim 1, wherein, if the first computing entity  
2 concludes that the first data entity is currently unavailable, the computing entity further  
3 writes an entry to a queue owner field in a queue to indicate an interest in accessing the  
4 first data entity.

1        12. (New)    The method of claim 11, wherein the computing entity also writes  
2 to a queue time field to indicate a period of time for which the entry to the queue owner  
3 field is valid.

1        13. (New)    The method of claim 1 further comprising reserving a disk on which  
2 the owner field and the time field are located to ensure exclusive access to the disk for  
3 the reading and writing of the owner field and the time field.

1        14. (New)    The method of claim 1, wherein, if a lease is obtained, the first  
2 computing entity also sets a renewal timer and, after the renewal timer expires, the first  
3 computing entity renews the lease by writing a new value to the time field.

1        15. (New)    A computer system comprising a first physical computer, a second  
2 physical computer, a data storage unit, a first data link for connecting the first physical  
3 computer to the data storage unit and a second data link for connecting the second  
4 physical computer to the data storage unit, the computer system further comprising:

5              a first virtual machine running on the first physical computer;  
6              a second virtual machine running on the second physical computer; and  
7              a file system stored on the data storage unit, the file system comprising:  
8                  a first data entity, the first data entity being usable by the first virtual  
9                  machine and by the second virtual machine; and  
10                 a lock for providing exclusive access to the first data entity, the lock  
11                 comprising an owner field and a time field, the owner field being used to  
12                 determine if the first data entity has been leased by a computing entity and

13 the time field being used to determine when a lease of the first data entity  
14 expires.

1               16. (New) The computer system of claim 15, wherein the first data entity is a  
2 file.

1           17. (New) The computer system of claim 16, wherein the file system further  
2 comprises a second file implementing a first virtual disk drive for use by the first virtual  
3 machine and a third file implementing a second virtual disk drive for use by the second  
4 virtual machine.

1               18. (New) The computer system of claim 17, wherein the first data entity is the  
2 second file.

1           19. (New) The computer system of claim 15, wherein the first data entity is a  
2 directory.

1           20. (New) The computer system of claim 15, wherein the file system further  
2 comprises a queue that may be used by a computing entity to indicate an interest in  
3 accessing the first data entity in the event that another computing entity has exclusive  
4 access to the first data entity.

1        21. (New) The computer system of claim 15, wherein the first data entity  
2 includes metadata and the lock is located in this metadata.

1           22. (New) The computer system of claim 15, wherein the first data link and the  
2 second data link are part of a data storage network.

1           23. (New) The computer system of claim 15, wherein the data storage unit  
2 comprises a disk drive.

1        24. (New) The computer system of claim 23 further comprising a disk  
2 reservation capability for providing exclusive access to the disk when accessing the  
3 lock.

1  
1        25. (New) The computer system of claim 15 wherein the first physical  
2 computer autonomously determines a first unique data value for identifying the first  
3 virtual machine in the owner field and the second physical computer autonomously  
4 determines a second unique data value for identifying the second virtual machine in the  
5 owner field.

1  
1        26. (New) The computer system of claim 15 wherein the first virtual machine  
2 is migrated from the first physical computer to the second physical computer and the  
3 first data entity remains usable by the first virtual machine and by the second virtual  
4 machine.

1  
1        27. (New) The computer system of claim 26 wherein the first data entity is a  
2 primary virtual disk drive for use by the first virtual machine.

1  
1        28. (New) The computer system of claim 15 wherein, if the first data entity has  
2 been leased for use by the first virtual machine and the first physical computer fails so  
3 that the lease on the first data entity cannot be released by the first physical computer,  
4 when the lease on the first data entity expires, the second physical computer is able to  
5 break the lease on the first data entity and begin using the first data entity.

1  
1        29. (New) The computer system of claim 28 wherein the first data entity is a  
2 virtual disk drive for the first virtual machine and, after the failure of the first physical  
3 computer and after the breaking of the lease on the first data entity, the first virtual  
4 machine is restarted on the second physical computer using the first data entity.

1       30. (New) A method for attempting to access a first data entity in a file system,  
2 the method being performed by a first computing entity, the file system also including  
3 one or more additional data entities that are concurrently accessible to at least one  
4 other computing entity, the file system including an owner field that can be used to  
5 determine whether the first data entity is in use by a computing entity, the method  
6 comprising:

7                  reading the owner field and determining whether the first data entity is in  
8 use by a computing entity;

9                  if the first data entity is not in use by a computing entity, writing to the  
10 owner field to take control of a lock on the first data entity; and

11                  if control of the lock is obtained, accessing the first data entity; or

12                  if control of the lock is not obtained, writing an entry to a queue owner field  
13 to indicate an interest in accessing the first data entity and waiting for an  
14 opportunity to access the first data entity.

1       31. (New) The method of claim 30 further comprising, if the first data entity is  
2 in use by a computing entity, reading a time field to determine whether a lease on the  
3 data entity has expired and, if the lease has expired, writing to the owner field to break  
4 the existing lease and to indicate an assumption of a new lease of the first data entity.

1       32. (New) The method of claim 31, wherein the first computing entity  
2 determines whether the lease has expired by reading a first value from the time field,  
3 delaying for a predetermined lease period and reading a second value from the time  
4 field, wherein the first computing entity determines that the lease has expired if the  
5 second value is the same as the first value, and the first computing entity determines  
6 that the lease has not expired if the second value is different from the first value.

1       33. (New) The method of claim 30 further comprising, if the first data entity is  
2 not in use by a computing entity, in addition to writing to the owner field to take control  
3 of the lock on the first data entity, writing to a time field to indicate when a lease of the  
4 first data entity expires.

1       34. (New) The method of claim 30, wherein the first data entity is a file.

1       35. (New) The method of claim 34, wherein the first data entity includes  
2 metadata and the owner field is located in this metadata.

1       36. (New) The method of claim 30, wherein the first data entity is a directory.

1       37. (New) The method of claim 30 further comprising reserving a disk on  
2 which the owner field is located to ensure exclusive access to the disk for the reading  
3 and writing of the owner field.

1       38. (New) The method of claim 30, wherein the first computing entity  
2 autonomously determines a data value that uniquely identifies the first computing entity  
3 and the first computing entity assumes a lock on the first data entity by writing the  
4 unique data value into the owner field.

1       39. (New) The method of claim 30 further comprising, if control of the lock is  
2 not obtained, in addition to writing an entry to a queue owner field to indicate an interest  
3 in accessing the first data entity, writing to a queue time field to indicate a period of time  
4 for which the entry to the queue owner field is valid.